

Taking Surveillance Out of the Shadows

Wiretapping and surveillance are shadowy subjects, covered as they are by the veils of secrecy that intelligence and law enforcement agencies reflexively apply to what they consider their most sensitive operations.

The subject raises fundamental and difficult questions regarding privacy, security, and accountability in a democratic society. Unfortunately, time and again, when the veils are lifted, we discover something perhaps even more unsettling than the specter of government abuse: the technologies at the root of surveillance systems are often deeply—sometimes embarrassingly—flawed.

Spying by Mistake

On 16 April, *The New York Times* reported that the US National Security Agency (NSA) had been inadvertently “over-collecting” purely domestic telephone and email traffic as part of its supposedly international warrantless wiretap program (E. Lichtblau and J. Risen, “Officials Say U.S. Wiretaps Exceeded Law,” p. A1). According to the article, the reasons for the over-collection errors were largely technical.

There’s a tendency to view wiretap systems, even controversial ones like the NSA’s, in strictly legal or political terms. Policy issues aside, we usually assume that the interception technology will correctly do whatever it’s supposed to do. But reality isn’t so simple. Although we don’t know all the details about how the NSA is car-

rying out international wiretaps in the US, what we do know suggests some questionable architectural choices that make the system especially vulnerable to collecting the wrong traffic by mistake.

Although the NSA wouldn’t discuss the technical details of its interception operation on the record, the *Times*’ sources said that the problem had to do with the agency’s inability to distinguish between domestic and international traffic in the data streams collected from US telecom networks. In particular, the little that has been disclosed about the NSA system includes collection components much farther inside the US telecom infrastructure than would be appropriate for intercepting the exclusively international traffic that the government says it wants.

According to court filings in the Electronic Frontier Foundation’s (EFF’s) lawsuit against AT&T, the NSA’s taps for international traffic are placed not, as we might expect, at the transoceanic cable landings that connect the US to foreign countries, but rather are inside switching centers that also handle a great deal of purely domestic traffic. Domestic calls are supposed to be excluded from the data stream the government re-

ceives, using specially configured network filtering devices the NSA has supplied to the carriers. But the taps are apparently in domestic backbone switches rather than, say, in the cable heads that leave the country, where international traffic is most concentrated (and segregated). This is, to say the least, a precarious way to ensure that only trans-border traffic will be collected, and an especially curious design choice given the NSA’s exclusively international mandate.

Compounding the inherent technical risks of the tap placement is another risk: the equipment panning for nuggets of international communication in the stream of (off-limits) domestic traffic is apparently composed entirely of hardware provided and configured by the government, rather than by the carriers. This is essentially equivalent to giving the NSA the keys to a phone company central office and hoping that they figure out which wires are the right ones to tap. We need not assume any malice on the NSA’s part to see how such a design invites error and risk.

As disturbing as the unauthorized surveillance might be, the sad fact is that domestic over-collection was a readily predictable consequence of the NSA’s poor design choices. Actually, it wasn’t just predictable—it was predicted. The technical community had been warning of this strange architecture’s risks since we first learned of it several years ago. In fact, we forewarned of precisely this over-collection scenario a year earlier in the Janu-

MATT BLAZE
University of Pennsylvania



ary/February 2008 issue of this publication, in an article entitled, “Risking Communications Security” that I coauthored with Steven Bellovin, Whitfield Diffie, Susan Landau, Peter Neumann, and Jennifer Rexford.

Déjà Vu

The NSA’s warrantless program wasn’t the first time that a large-scale wiretap architecture has been abused or compromised in ways that the designers didn’t anticipate. In 2005, officials discovered that more than 100 cell phones in Athens, Greece, mostly belonging to politicians (ranging from the local mayor to the prime minister), were being illegally wiretapped. (Vassilis Prevelakis and Diomidis Spinellis published a fascinating technical analysis of the case in the July 2007 issue of *IEEE Spectrum*.) We still don’t know who did it or why, but we do know how: the culprit exploited weaknesses in the special “lawful interception” features built into the Greek telephone infrastructure. These interfaces were intended to make it easier for the police to monitor sus-

pected criminals’ calls, but, in this case, the criminals used them to do the monitoring. Of concern to non-Hellenic readers is the fact that 1994’s Communications Assistance for Law Enforcement Act (CALEA) mandated the same mechanisms for inclusion in US telecom systems as were abused in Greece.

CALEA, like the NSA wiretap program, was (and remains) controversial, but, again, the public controversy focused largely on the philosophical rather than the technical. In CALEA’s case, the debate focused on the legal, economic, and social implications of requiring the communications systems that millions of law-abiding people use to be universally “wiretap-ready” for the relatively rare (in the US, at least) case where law enforcement needs to monitor someone. These questions, as important as they are, pale next to an even more difficult problem—how, exactly, can we make our entire communications infrastructure susceptible to tapping by the good guys without also making it more vulnerable to abuse by the bad guys? The hard technical

reality, as we saw in Greece, is that we probably can’t—whatever our policy intentions might be.

As bad as defective eavesdropping systems might be from a privacy perspective, the problems can be even worse from the perspective of the authorized wiretappers—the law enforcement and intelligence agencies that rely on surveillance for reliable evidence about their targets. Almost every secretly developed wiretapping system about which details have emerged has turned out to suffer from flaws that make it possible for a target to evade surveillance or compromise collected evidence.

In fact, wiretapping technology’s recent history is absolutely littered with systems that turn out to be unable to effectively wiretap. A famous example was the NSA’s ill-fated Clipper chip, released in 1993, which had a weakness (discovered a year later after some reverse-engineering) that made it easy for criminal users to defeat the key escrow features—the whole purpose of the system—while still using the cipher algorithm. But that’s only the most prominent failure. A more recent example came to light in 2005, when Micah Sherr, Eric Cronin, Sandy Clark, and I discovered that the “loop-extender” taps that law enforcement has used for years to monitor analog telephone calls employed vulnerable in-band signaling that lets criminals remotely disable the tap and introduce false data into the logs, just by sending a few audio tones down their phone lines. CALEA systems, used today to tap cell phones and data services, suffer from similar weaknesses that can neutralize their effectiveness against criminals.

Watching the Watchers

Wiretapping modern communications systems correctly and securely is an extremely complex, subtle, and difficult problem. But although computer networks and

telephone switches are carefully engineered and tested to work under a wide range of conditions before they're deployed (and quickly fixed when problems are discovered), the hardware and software systems law enforcement and security agencies use to tap them usually aren't. Instead, wire-tapping systems are typically built in much the same way as they're used: in the shadows, away from public view or careful engineering scrutiny. And the results, as we've seen time and again, are exactly what you would expect.

In other words, surveillance systems aren't exempt from the same basic engineering reality that affects computing generally—making complex systems secure is just plain hard. No easy recipe exists for building them, but the most useful thing years of software engineering research has taught us

is that security comes only from extensive, wide scrutiny. (This is why widely used open source software often enjoys a security advantage, despite the fact that the attackers get to look at the source code along with the defenders.)

This might not be a pleasant answer from the perspective of those who build or rely on wiretap systems because it suggests that the secrecy that traditionally shrouds their design is effectively hindering their usefulness.

Fortunately, there's no reason to hide the design. Although we might need to keep secret *who* is being targeted, there's no benefit here to keeping secret *how*—in a properly designed system, knowing the mechanism should be of no help to criminals trying to evade it. And as we openly debate wiretap

policy, we must start to give equal attention to the technology that implements it, lest we read more headlines about failed surveillance systems that spy on the wrong people. Ultimately, if surveillance is to properly protect us, it has to come out of the shadows. □

Matt Blaze directs the Distributed Systems Laboratory at the University of Pennsylvania. His research focuses on computer and network security, cryptography, and the relationship of technology to public policy. Blaze has studied surveillance technology since 1994, when he discovered basic design flaws in the NSA's Clipper chip. He has a PhD in computer science from Princeton University. Contact him at blaze@cis.upenn.edu.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

RUNNING IN CIRCLES LOOKING FOR A GREAT COMPUTER JOB OR HIRE?

The IEEE Computer Society Career Center is the best niche employment source for computer science and engineering jobs, with hundreds of jobs viewed by thousands of the finest scientists each month - **in *Computer* magazine and/or online!**

 careers.computer.org
<http://careers.computer.org>

The IEEE Computer Society Career Center is part of the *Physics Today* Career Network, a niche job board network for the physical sciences and engineering disciplines. Jobs and resumes are shared with four partner jobboards - *Physics Today* Jobs and the American Association of Physics Teachers (AAPT), American Physical Society (APS), and AVS: Science and Technology of Materials, Interfaces, and Processing Career Centers.

 IEEE
computer
society

- > Software Engineer
- > Member of Technical Staff
- > Computer Scientist
- > Dean/Professor/Instructor
- > Postdoctoral Researcher
- > Design Engineer
- > Consultant

